

# **SANDIA REPORT**

SAND2012-8437

Unlimited Release

Printed October 2012

## **An approach for estimating the uncertainty in ParaDiS predictions**

J. Ray, H. N. Najm, M. Rhee and A. Arsenlis  
Sandia National Laboratories, CA  
Lawrence Livermore National Laboratory, CA

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energys National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



**Sandia National Laboratories**

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from  
U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
Online ordering: <http://www.doe.gov/bridge>

Available to the public from  
U.S. Department of Commerce  
National Technical Information Service  
5285 Port Royal Rd  
Springfield, VA 22161

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: [orders@ntis.fedworld.gov](mailto:orders@ntis.fedworld.gov)  
Online ordering: <http://www.ntis.gov/ordering.htm>



# **An approach for estimating the uncertainty in ParaDiS predictions**

Jaideep Ray and Habib N. Najm  
Sandia National Laboratories, Livermore CA 94551-0969  
{jairay,hnnajm}@sandia.gov

Moono Rhee and Athanasios Arsenlis  
Lawrence Livermore National Laboratory, Livermore CA 94551-0808  
{rhee1,arsenlis1}@llnl.gov

## **Abstract**

This report outlines an approach for computing the uncertainties in the predictions of computationally expensive models. While general, we use ParaDiS, a dislocation dynamics simulator originating in Lawrence Livermore National Laboratory, as the target application. ParaDiS is a mesoscale model, and uses submodels constructed/upscaled from microscale (molecular statics and dynamics) simulations. ParaDiS outputs, in turn, are upscaled and used in continuum (macroscale) simulations, e.g., those performed by ALE3D. This report addresses how one may quantify the uncertainties introduced by upscaling (both from microscale to mesoscale, and mesoscale to continuum), and the dependence of uncertainties in ParaDiS predictions on those of the inputs. This dependence is established via sensitivity analysis, and we address how this may be performed with a minimum of ParaDiS runs, given its immense computational cost. This includes constructing a smaller version of the model, sparse sampling of the parameter space, and exploiting the asymptotic nature of the time-evolution of the model outputs. The report concludes with a discussion of the computational resources required to perform this uncertainty quantification study.

## **Acknowledgment**

This work was supported by Sandia National Laboratories' Advanced Simulations and Computing (ASC) program and the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U. S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Contents

1	Introduction	7
2	Literature review	11
2.1	Bayesian calibration and model selection	11
2.2	Surrogate models	12
2.3	Sparse sampling	13
3	Description of the model	15
3.1	The inputs	16
3.2	Outputs	17
3.3	The numerical method	17
4	Outline of the uncertainty quantification study	19
4.1	Uncertainties in Type I parameters	19
4.2	Uncertainties in Type II parameters	19
4.3	Uncertainties in Type III parameters	20
4.4	Designing a computationally inexpensive ParaDiS problem	30
4.5	Sampling strategy for ParaDiS	30
4.6	Estimating model uncertainty in $\mathcal{M}_2$	31
4.7	Execution plan and deliverables	33
5	Computational issues and resources	38
6	Summary	40
	References	44

# Figures

1	Fig. 3 from [1]. Blue and green symbols were read off with a graph paper and used as MD “observations” in the ParaDiS input characterization.	20
2	Marginalized posterior distribution of $\beta_0$ (top) and $\beta_T$ (bottom).	22
3	Posterior predictive tests for 300K (top) and 700K (bottom), when 2 parameters are left uncertain. The tests do not match the variability in the data. The solid line is the median from the posterior predictive tests, and the error bars show the 99% credibility interval.	23
4	Pairwise plot of the 6 parameters being estimated.	24
5	Prior and posterior distributions for $\beta_0, \beta_T$ , (top) $\alpha_0, \alpha_1$ , (middle) $\chi_0$ and $\chi_1$ (bottom). The nominal values, currently used in ParaDiS, are plotted as vertical lines. Red lines denote the distributions computed with just $(\beta_0, \beta_T)$ assumed uncertain (Sec. 4.3.1).	26
6	Posterior predictive tests for 300K (top) and 700K (bottom). We see that we capture the variability of the MD data. The solid line plots the median of the posterior predictive test and the error bars cover the 99% credibility interval.	27
7	Posterior for $\sigma$ , using 6 uncertain parameters.	28

8	Verification rank histograms for the 2-parameter model (top) and the 6-parameter model (bottom). Thirty observations are binned into 30 equally-spaced bins. We clearly see that the calibration of the 2-parameter model is under-dispersive, whereas the 6-parameter model displays a good calibration.	29
9	Data for fitting Eq. 2 and the resulting fit. Top: ParaDiS output in black (plotted against the left Y-axis) and the stress-strain curve (blue; plotted against the right Y axis), showing the plastic region. The red region is the plastic-deformation data used for fitting Eq. 2. The green line is the fitted model, using optimal values of $R, \rho_{sat}$ . Bottom left: $(R, \rho_{sat})$ samples from the fitted distribution. We see that there is very little uncertainty in the fit of $\mathcal{M}_2$ ; the distribution is very narrow, but $(R, \rho_{sat})$ clearly assume correlated values. Bottom right: A zoom of the fitted (very narrow) Gaussian. . . . .	34
10	Predictions with the fitted model. Top: Using values of $(R, \rho_{sat})$ sampled from the distribution plotted in Fig. 9 (bottom), we prediction an ensemble of 100 $\dot{\rho}$ evolutions, plotted in green. The ParaDiS output is plotted in black. Bottom: Distribution of predictions $\rho(t)$ (arising from Eq. 2 and $(R, \rho_{sat}) \sim \text{MVN}(\mu, \Gamma)$ ) at 6 different points in time (each corresponding to a different color). The dashed lines are the ParaDiS results. The black plot is at the earliest time, the bold magenta at the latest time. . . . .	35

## Tables

1	The number of Clenshaw-Curtis sparse grid quadrature points for various levels and dimensionalities. As a comparison, a product grid with only 4 points per dimension in 10-d consists of $4^{10} > 1000000$ points. . . . .	14
2	Mean, standard deviation and bounds of independent truncated normal prior distributions for the parameters $\mathbf{p}$ . . . . .	21

# An approach for estimating the uncertainty in ParaDiS predictions

## 1 Introduction

This report outlines an approach by which one can quantify the uncertainty in the predictions of a computationally expensive model, due to uncertainties in its inputs. The model is also high-dimensional in its input parameter-space, and some of its inputs are expected to be correlated, if aphysical situations (which the model cannot address) are to be avoided. In this report, we describe how the uncertainties in the input parameters may be modeled as well as a method for constructing a computationally inexpensive surrogate model for performing the uncertainty quantification (UQ) study.

The model in question is ParaDiS [2], a simulation code for dislocation dynamics (DD). It is currently being developed by, among other institutions, Lawrence Livermore National Laboratory (LLNL). The UQ study is complicated by the manner in which ParaDiS is used - it serves as a mesoscale model in a multiscale upscaling chain, whose final product (an upscaled model) is used as a constitutive submodel in continuum simulations. In this multiscale upscaling chain, molecular dynamics (MD) simulations are used to create a model for the stress  $\sigma_e$  required to move an isolated crystal dislocation with velocity  $v$ , at a given temperature  $T$ . The MD data is fitted with an algebraic model, which we will refer to as the first upscaling model,  $\mathcal{M}_1$ . The parameters of  $\mathcal{M}_1$  can be best described as random variables with a distribution characteristic of the fit of  $\mathcal{M}_1$  to MD data. DD models the interaction of dislocations in a single crystal, including creation and annihilation of dislocations, leading to high-strain rate behaviors as well as strain hardening. ParaDiS outputs consist of the evolution of dislocation density over time under an applied stress (or strain rate). The dislocation density evolution is, in turn, modeled and used as a constitutive model in continuum simulations (e.g., in ALE3D [3]). We call this the second upscaling model,  $\mathcal{M}_2$ . The upscaling itself introduces uncertainties, *even if ParaDiS input were known with absolute certainty*. An example of ParaDiS being used as a part of a multiscale chain of tools can be found in [1]. DD thus acts as a bridge between MD and the continuum scale.

In our work, we will focus on the final predictive uncertainties in  $\mathcal{M}_2$ , due to uncertainties in the parameters of  $\mathcal{M}_1$ , the shortcomings of  $\mathcal{M}_2$  itself as a representation of ParaDiS outputs, as well as artifacts introduced into ParaDiS simulations due to numerical errors/approximations. We will focus solely on tantalum (Ta) as the material of choice. We will assume that the upscaling tool-chain and the use-case is identical to the description in [1].

The inputs to ParaDiS can be divided into four classes:

- Type I parameters: Conventional parameters, that have a nominal value and an uncertainty about them. ParaDiS has 8 such parameters (`pois`, `burgMag`, `YoungModulus`, `rc`, `MobClimb`, `MobLine`, `MobGlide`, `MobPieirls`). These can be modeled as being independently distributed.
- Type II parameters: These are “operational” parameters, i.e. parameters that are varied to obtain ParaDiS outputs over a space where (a) ParaDiS is applicable and (2) where ALE3D needs its constitutive law to be evaluated. The 3 parameters are  $T$ , the temperature,  $\dot{\epsilon}$ , the strain rate and  $G$ , the shear modulus, being a proxy for pressure.
- Type III parameters: ParaDiS uses a constitutive law (the upscaling model  $\mathcal{M}_1$ )  $\sigma_e = f(v, T; \vec{p})$  in it, which relates the stress required to move an isolated dislocation, with velocity  $v$ , at a temperature  $T$ .  $\vec{p}$  are the parameters of  $\mathcal{M}_1$ . This model is obtained by fitting to MD data. Due to shortcomings of the functional form of the model, and jitter in the MD data, there is considerable uncertainty in the values of the model parameters. Typically, one represents the MD data as  $\sigma_e^{data} = f(v, T; \vec{p}) + \eta$ , where the uncertainties on  $(\vec{p}, \eta)$  have to be obtained.  $\eta$ , in turn, has to be modeled (e.g.,  $\eta \sim N(0, \sigma_\eta^2)$ ), but a joint distribution on  $(\vec{p}, \sigma_\eta)$  could be obtained. This joint distribution has to be honored when performing a parametric UQ study on ParaDiS i.e., these parameters should not be modeled as independently distributed. There can be up to 9 such parameters (a more probable number is 6).
- Type IV parameters: Numerical parameter, governing convergence, time-step, order of FM expansions etc. Ideally, ParaDiS outputs should be independent of numerical settings (convergence with respect to grid etc, for a proper run), but we will have to estimate these sensitivities, for reasons discussed in Sec. 4.4. These are `maxDT`, `deltaTT`, `nextDT`, `dtIncrementFact`, `dtDecrementFact`, `dtExponent`, `dtVariableAdjustment`, `rTol`, `rmax`, `splitMultinodeFreq`, `fmMPOrder`, `fmTaylorOrder` - 12 in all.

ParaDiS has many outputs (predictions); however, only the following are of interest to us:

- *Dislocation density versus time.* This is a function of the strain rate  $\dot{\epsilon}$ . This output is fitted with an upscaling model ( $\mathcal{M}_2$ ; Eq. 11 in [1]), and used in continuum simulations. The upscaling model is fitted only when the metal begins to display plastic behavior. We are interested in quantifying the uncertainty in the fitted-model’s predictions, due to shortcomings of the fit to data from a particular ParaDiS run *and* the uncertainty in ParaDiS outputs, due to uncertainties in its inputs.
- *The stress-strain relationship.* This is used to identify the plastic region, prior to fitting the upscaling model described above.

In this report, we will address the following issues:

- *Examine the distribution of Type III parameters:* We will examine distribution of  $\vec{p}$  in the dislocation velocity model  $\sigma_e = f(v, T; \vec{p})$ , by fitting to MD data. Typically, one obtains point estimates of these parameters. Instead, we will focus on the identifiability of the six parameters that constitute  $\vec{p}$  i.e., whether the posterior distributions of the parameters, after fitting the model to MD data, differ appreciably from the prior. Those parameters that cannot be identified (i.e., the MD simulations are insensitive to their variation) can be safely set to a constant and removed from any UQ study.
- *Construct a “mini” version of ParaDiS, that can be run quickly:* A typical ParaDiS run takes  $O(10^7)$  CPU-hours; thus we cannot use ParaDiS as-is in our UQ study. We seek to construct a “mini” version that can be completed in approximately  $O(10^5)$  CPU-hours, so that it may be run repeatedly within the context of our study. We seek to achieve this by reducing the domain size as well as curtailing the number of timesteps run. However, this number has to be large enough such that (1) the numerical experiment displays plastic behavior (a strain of about 2–5%) and (2) the evolution of dislocation density over time attains the asymptotic behavior characteristic of saturation. The last is required so that we may fit the upscaling model  $\mathcal{M}_2$ .

However, construction of a “mini”-ParaDiS will lead to numerical errors and consequently the effect of Type IV parameters on the model outputs will have to be characterized.

- *Sampling design:* We will outline how one may sample the  $(8 + 3 + 9 = 20)$  dimensional space (Type I, II and III parameters) using “mini”-ParaDiS. We expect that the examination of the identifiability of Type III parameters may yield a modest reduction in parametric dimensionality, but the bulk of the parametric study will involve sparse sampling. Which methods may yield a successful outcome will be discussed and shortlisted in this report.
- *Upscaling model  $\mathcal{M}_2$ :* The upscaling model  $\mathcal{M}_2$  (Eq. 11 from [1]) provides a functional form that can be fitted to data. However, the discrepancy between ParaDiS and  $\mathcal{M}_2$  cannot be modeled as noise (ParaDiS outputs are noiseless) and we will identify a probabilistic modification of  $\mathcal{M}_2$  that serves to bridge the discrepancy.

In Section 4, the report will present preliminary results on three of the issues mentioned above, indicating that their solution is eminently feasible. For the fourth component, sampling design, will be reviewed (in Section 2); however, no preliminary results will be presented. They form major components of an UQ study of ParaDiS. We will end this report with an outline how these components can be integrated to execute an UQ of ParaDiS.

This page intentionally left blank

## 2 Literature review

As mentioned in Sec. 1, the key to performing an UQ study on ParaDiS is to estimate the uncertainties in Type III parameters (which govern  $\mathcal{M}_1$ ) as well as the uncertainties incurred when ParaDiS results are upscaled/fit using  $\mathcal{M}_2$ . These uncertainties due to model fitting can be assessed if one uses Bayesian calibration. In this section, we review some of the methods, and some past work that may be useful for our purposes.

UQ analysis of ParaDiS, which will require multiple model evaluations, is intractable using ParaDiS as-is; instead, we will take recourse to a computationally inexpensive surrogate. A surrogate model is created by sampling the parameter space sparsely, and “fitting” a functional form. We will also review sampling methods and surrogate models in this section.

An UQ approach that allows one to map uncertainties in input parameters to uncertainties in model prediction, via a surrogate model, essentially creates a probabilistic (surrogate) simulator. This approach can be used to test the change in ParaDiS’s predictive skill if its submodels e.g.  $\mathcal{M}_1$  were to be replaced by a competing one. The new submodel would have different parameters with different levels of uncertainties. Being able to select between competing submodels would be a *use* of the approach being described here. We will review how model selection may be performed in a probabilistic manner, given experimental data.

### 2.1 Bayesian calibration and model selection

Model calibration is commonly thought of as the manipulation of model parameters so that model predictions may match measurements (from the field or laboratory experiments). Typically this is done via an optimization procedure, where a cost function (usually the  $L_2$  norm of the data - model mismatch) is minimized. In Bayesian methods, this mismatch is modeled (often by assuming that individual mismatches are uncorrelated and furthermore, are distributed per a Gaussian). Thus, given a proposed set of parameter of values, one can calculate the mismatch; then, per the error distribution model, one can calculate the likelihood of observing that error, and by extension, the chosen set of parameters. If there are any prior beliefs regarding the choice of parameters (a probability distribution over the parameter space), the likelihood can be modified accordingly. Per Bayes theorem, this is proportional to the *posterior* distribution on the parameters, conditional on the observations. For more information on such a statistical/probabilistic formulation of a parameter-estimation problems (or Bayesian inverse problems), see [4].

The posterior distribution cannot be evaluated in a naïve manner e.g., by gridding up the parameter space (which in our case is 20-dimensional). Instead one samples from it (obtaining samples from the parameter space) and bins (histograms) them to illustrate the geometry of the posterior distribution. The sampling is performed using Markov chain Monte Carlo (MCMC) methods which are described in [5]. In particular, for our work, we will use an adaptive MCMC method called DRAM [6]. Adaptive MCMC reduces the number of

samples required to obtain a “converged” a posterior distribution, especially if it has a complex, contorted geometry. The samples drawn from a posterior distribution can be used to construct a response surface model of the posterior distribution. Fitting high-dimensional response surfaces to a cloud of points will be discussed in Section 2.2 within the context of surrogate models.

One of the uses of an approach to generate probabilistic surrogate models for ParaDiS is to be able to discriminate between competing submodels e.g., for dislocation mobility ( $\mathcal{M}_1$ ), that might be implemented in ParaDiS. The competing submodels will have different parameters, with different posterior distributions. These, in turn, will lead to a corresponding distribution of predictions of dislocation density evolution. In the event laboratory data is available, the question arises how one may discriminate between the competing submodels.

For this purpose, we borrow ideas from Bayesian model selection. Given samples from the posterior distribution, one may simply calculate the Deviance Information Criterion for the competing versions of ParaDiS [7]. Alternatively, one may calculate various goodness-of-fit metrics e.g. cumulative rank predictive score, that are defined on the predictive skill of an ensemble of predictions with respect to observations. Their definitions and implications can be found in [8, 9]. Finally, samples from the a distribution of ParaDiS parameters can be used to discriminate between competing version of ParaDiS (if experimental observations are available) using Bayes factors [10]. The surrogate model for ParaDiS will be required for the high-dimensional integrations involved in the calculation of Bayes factors.

## 2.2 Surrogate models

Surrogate models are computationally inexpensive analogs of expensive computational models. These surrogates approximate one (and sometimes more) outputs of a model as a function of model inputs. The surrogate models, sometimes also called response-surface models, typically do not have any scientific/phenomenological arguments underlying their construction and can be likened to “curve-fitting”. The primary issues involved during surrogate model construction are (1) minimizing the number of expensive-model evaluations to generate the data to which surrogates are fit (generally accomplished with some kind of sampling) and (2) minimizing the difference between predictions/outputs of the expensive model and its inexpensive surrogate. Descriptions of the issues involved in generating surrogates can be found in [11, 12, 13, 14]. Surrogate models are popular in inversion and optimization studies since those studies involve repeated evaluation of models for different parameter values.

One of the ways of constructing a surrogate for ParaDiS is to use polynomial chaos (PC) expansions to model its outputs, as a function of its inputs. The polynomial chaos (PC) was defined first by Wiener [15], and it since has found a significant number of applications in a number of engineering fields [16, 17, 18, 19, 20]. This approach consists of approximating a generic random variable in terms of standard random variables through a spectral polynomial expansion. In the context of this report, dislocation density evolution

(or proxies thereof) will be cast as random variables that are functions of uncertain input parameters. These spectral approximations can be constructed using a relatively small number of ParaDiS evaluations, and can accurately represent smooth input-output dependencies. If ParaDiS is seen to exhibit non-smooth behavior, several domain partitioning methods are available [21, 22, 23]. This generates a series of sub-domains where models have a smooth behavior, thus enabling the use of efficient spectral approximations in each region. Alternatively, if ParaDiS outputs were to display smooth, near-stationary behavior, with few localized areas of large gradients, surrogate model methods based on Gaussian processes (or treed Gaussian processes) could be used to create surrogates [12, 13, 24, 25].

Marzouk *et al.* [26] proposed using surrogate models based on PC expansions in order to accelerate Bayesian inferences. This approach was followed by several authors in a wide range of scientific fields; for source and parameter estimation in porous media [27, 28], analysis of supersonic combustion [29], and stochastic data assimilation [30]. A use of PC surrogates in global sensitivity analysis is in [31].

## 2.3 Sparse sampling

If a one-dimensional function  $f(x)$  needs to be integrated, it is sufficient to evaluate it at  $Q$  quadrature points; this will yield an exact answer for a polynomial  $f(x)$  of order  $2Q - 1$  or less (for Gaussian quadratures). The set of quadrature points forms a 1D grid; trapezoidal or Simpson's rules are quadrature grids.  $p = 2Q - 1$  is referred to as the *precision* of the grid. In case the function is extended to higher dimensions e.g.  $f(x_1, x_2, \dots, x_d)$ , the 1D grid can be naïvely replicated in  $d$  dimensions, leading to  $Q^d$  grid-points. Within the context of an expensive  $f(\mathbf{x})$ , evaluating the function at each grid point becomes intractable. Sparse Clenshaw-Curtis (CC) grids and other sparse quadrature methods become relevant in such a context.

Sparse grids date back to Smolyak [32], whose work serves as a starting point for efficient methods of high-dimensional quadrature construction based on uni-dimensional rules [33, 34, 35]. For details of the construction, see [36, 37]. Sparse grids are composed of *levels*, with each level  $L$  consisting of a set of points unevenly distributed in space. The spacing between points is decided by a 1D formula. If one desires a higher precision, a new level, with a larger number of points is added;  $p = 2L - 1$ . Table 1 shows the growth of precision  $p$  and number of grid-points  $N$  as more levels are added. We see that a  $L = 5, d = 10$  CC grid contains almost 9,000 points, and has a precision of 9. Note that the highest-order polynomial that can be integrated exactly on a  $d$ -dimensional,  $L$  level grid is  $x_1^{q_1} x_2^{q_2} \dots x_d^{q_d}$ , where  $\sum_{k=1}^d q_k = p = 2L - 1$ . For an  $L = 5, d = 10$ , which would be clearly intractable with ParaDiS, we would not be able to go beyond  $5^{th}$  order terms. Since we expect terms of the form  $x_i x_j x_k \dots$ , it is thus clear that a significant reduction of the parameter space will be required to perform an UQ study of ParaDiS if its outputs display complex/high-order dependencies on input parameters.

**Table 1.** The number of Clenshaw-Curtis sparse grid quadrature points for various levels and dimensionalities. As a comparison, a product grid with only 4 points per dimension in 10-d consists of  $4^{10} > 1000000$  points.

Level $L$	Precision $p = 2L - 1$	$N$ ( $d = 2$ )	$N$ ( $d = 5$ )	$N$ ( $d = 10$ )	$N(d)$ General
1	1	1	1	1	1
2	3	5	11	21	$1 + 2d$
3	5	13	61	221	$1 + 2d + 2d^2$
4	7	29	241	1581	$1 + \frac{14}{3}d + 2d^2 + \frac{4}{3}d^3$
5	9	65	801	8801	$1 + \frac{20}{3}d + \frac{22}{3}d^2 + \frac{4}{3}d^3 + \frac{2}{3}d^4$
6	11	145	2433	41265	...
7	13	321	6993	171425	...
8	15	705	19313	652065	...

### 3 Description of the model

ParaDiS is a computational model, being developed in LLNL, that computes the interaction of crystal dislocations with each other when a crystal of metal is subjected to a strain-rate. Pressure and temperature, during a ParaDiS run, are held constant. Fundamentally, as the metal deforms elastically, the crystal planes move reversibly. As the material passes into plasticity, dislocations are born and the deformation becomes irreversible (and the strength drops). Both screw and edge dislocations are modeled. Screw dislocations are modeled in detail and the contributions to the model outputs, due to edge dislocations, is simply taken as a fraction of the contributions by screws [1]. As the imposed strain rate continues to deform the material, the dislocation density increases and the dislocations themselves start interacting with each other and resisting the strain i.e., a larger stress is imposed to keep the strain-rate constant. This increase in material strength is called strain-hardening and plays an important part in the the functioning of metal devices under extreme pressures and temperatures.

The interaction of dislocations with each other is facilitated by the stress field, and the movement of dislocations is far slower than the rate at which the stress field propagates in the material (i.e., it is an elliptic problem). The dislocation evolution problem is addressed using a Fast Multipole Method (FMM). Details of the numerical method are in [38]. Problems of physical interest tend to be enormous and it is the model's weak scalability that is exercised in parallel computing. The weak-scalability is quite impressive.

The kind of physics simulated by ParaDiS – dislocation dynamics – functions at the mesoscale level. It is far coarser than molecular statics and dynamics, and expects certain constitutive models to be computed *a priori* from such simulations. These molecular simulations address the evolution (movement/velocity) of a single dislocation under various degrees of applied stress, at a fixed (extreme) pressure and temperature. DD, in turn, simulates processes that involve interactions between dislocations and reproduces the stress–strain relationship in a solid, as well as the temporal evolution of dislocation density and dislocation flux under an applied strain rate (the former rises to an asymptotic limit; the latter displays a rise and a fall, before settling to a constant). Thus DD forms a step in a multiscale chain of models, stretching from molecular statics (used to calculate activation energies of dislocations), to dislocation mobilities as a function of stress (molecular dynamics), which then feeds these “constitutive” relations into DD. The output of DD, properly upscaled, informs continuum models. This entire chain, along with illustrative examples, has been discussed in [1].

The ParaDiS code is freely downloadable [2], however, most modifications of DoE ASC (Advanced Simulation and Computing) interest e.g., models specific to extreme pressures, temperatures and relatively exotic metals (tantalum, vanadium etc) are only available in the LLNL version of the model. The model has not been cleared for public release.

ParaDiS simulates DD in a 3D box. Typically, the box-size is 5 – 10 $\mu$ . Each simulation was run on 4,000–16,000 CPUS on the Blue Gene/L, and simulations required to meet

an ASC L2 milestone consumed about 200 million CPU-hours, primarily because of the large strains (i.e large number of timesteps) that had to be achieved. Consequently, such a problem configuration is not really suitable for repeated runs, as would be needed for uncertainty quantification. The restriction comes from the number of dislocations that arise under an imposed strain rate, which requires one to restart the problem on a larger number of processors. Also depending on the model used to relate dislocation velocity to imposed (local) stress, the timestep can drop to around  $10^{-12}$  seconds, requiring  $O(10^5)$  timesteps to calculate physically useful outputs. Smaller problems ( $2\mu$  box-size) are possible but show oscillations in model output (stress-strain curve, dislocation density and flux as a function of time etc) as the perturbations circulate in the domain (under periodic boundary conditions), slowly attenuating by energy dissipation due to friction and PdV work. Increasing the domain size allows the energy dissipation to occur over a larger volume and decreases its impact on dislocations. All ParaDiS outputs are averages over the box-volume; thus averages over a larger number of dislocations also result in a cancellation of fluctuations and a smoother output.

### 3.1 The inputs

ParaDiS obtains inputs on dislocation mobility (velocity of a dislocation as a function of the local stress) from molecular dynamics simulations. The model is formulated as

$$\sigma_e = M(\tau^*(v, \theta)\tau_p(p) + \hat{\tau}(\rho, p) + \tau_a)$$

where  $v$  is the dislocation velocity,  $\theta$  is the temperature in Kelvins,  $p$  is the pressure (above atmospheric pressure),  $\rho$  is the dislocation density and  $\tau_a$  is a constant, which is small for certain materials (e.g., vanadium and tantalum) as it is related to the grain structure. This equation is Eq. 2 in [1].

The Pierls stress  $\tau_p$  is the strength of the material at 0 K. Pierls stress at a reference pressure is obtained from molecular statics simulation for a number of (low) shear moduli  $G$ . A form  $\tau_p = \tau_{p,0}(G/G_0)$  is used to scale  $\tau_p$  to different shear moduli. Since  $G$  is a function of pressure, scaling to higher pressure is simply done using Eq. 4 in [1]. The two empirical scalings  $\tau_p = \tau_{p,0}(G/G_0)$  and Eq. 4 are *one (but a minor) source of uncertainty* for vanadium and tantalum. They are routinely used in all branches of material science.

The second expression  $\hat{\tau}(\rho, p) = \beta b G \sqrt{\rho}$ , where  $\beta$  is a constant of order 1,  $b$  is the Burgers vector,  $G$ , the shear modulus is pressure dependent, and  $\rho$  is the dislocation density. Dislocation density changes in time but saturates at a value dependent on the strain-rate  $\dot{\epsilon}$  imposed during a ParaDiS run. This too is a *minor source of uncertainty* in most metals.

*The largest source of uncertainty is  $\tau^*(v, \theta)$ .* It models the work-hardening of a metal due to large deformations and is estimated from molecular dynamics (MD) simulations. At low values of  $v$ ,  $\tau^*$  shows complex dependence on  $\theta$  for vanadium and none for tantalum. At high  $\theta$ , the behavior of  $\tau^*$  is linear with  $\theta$  (for tantalum) and approximately linear for vanadium (this region is called the phonon drag region). However, for practical purposes,

fidelity is desired in regions where phonon drag is negligible. This is where complex temperature dependence on  $\theta$  is observed and where MD simulations are performed.

In MD simulations, one initializes a dislocation in a box and imposes a shear stress  $\sigma_e$  on it. The position of the dislocation is measured as a function of time and its velocity estimated by fitting a slope. After initial transients, a slope is observed (see Fig. 6 in [39] and Fig. 8 in [40]). MD runs are repeated at different  $\sigma_e$ , with the pressure and temperature held constant. The numerical experiments are then repeated at multiple temperatures. Feasible MD simulations are typically inaccurate below 300K and consequently the runs are usually performed for the range 200K – 1100K. Also, in the absence of any other dislocations, the strain-rates observed in MD are about 5-orders of magnitude larger than what ParaDiS predicts.

The  $(v, \sigma_e)$  pairs obtained from MD simulations are fitted using a functional form (Eq.5, 6, 7, 8 in [1]). This functional form has 8 parameters *and is the primary source of uncertainty*. However, of these, 3 are related to the phonon drag regime, which is not of much practical use. However 4 parameters, representing  $\sigma_e - v$  variation in the temperature-dependent (non-phonon-drag) region, are critical and the uncertainty in their estimates, to first order, are most desired. These uncertainties arise from two sources (1) shortcomings of the functional form itself and (2) noise in the MD data. An MD run, repeated multiple times, for the same pressure, temperature and  $\sigma_e$  settings will provide multiple (but close) values of  $v$ . The sensitivity of ParaDiS outputs to these 4 parameters is critical. If the uncertainty in ParaDiS predictions due to shortcomings of the functional form (and noise in MD simulations) are large, then perhaps MD simulations need not be run to very large times or with large box-sizes (to obtain “converged” values of  $v$ ).

## 3.2 Outputs

ParaDiS provides 3 outputs, which are upscaled for use in continuum codes. These are (1) the stress-strain ( $\sigma - \epsilon$ ) relationship (2) the expression for the time evolution dislocation density and (3) dislocation flux evolution. The expression for the time evolution of dislocation density is

$$\dot{\rho} = R \left( 1 - \frac{\rho}{\rho_{sat}(\dot{\epsilon})} \right) \dot{\epsilon}.$$

$\rho_{sat}$  is obtained by running multiple ParaDiS runs at different values of  $\dot{\epsilon}$  and is fitted by a  $\rho_{sat} \sim \dot{\epsilon}^a$  expression (Fig. 4 in [1]). Both  $R$  and  $\rho_{sat}(\dot{\epsilon})$  are *sources of output uncertainty*.

## 3.3 The numerical method

The numerical method is somewhat less relevant to the non-intrusive, sampling-based uncertainty quantification study contemplated here and is provided to provide context for

some of the Type IV parameters discussed in Sec. 4. Full details are in [38] and only a summary is reproduced here.

In a single crystal, dislocations are often modeled as curved lines which may intersect others. Numerically, they are modeled as straight-line segments, broken at nodes to accommodate curvature, or to represent the intersection of dislocations. Nodes, and segments connecting them, constitute a dislocation network. In ParaDiS, nodes act as sites which experience force and are subjected to dynamics; the segments connecting nodes can extend or shrink (subject to bounds on their length). Nodes can be born or merged/annihilated. Force on a node is calculated as a local spatial derivative of the energy field; the field itself is a sum of energy due to long-range elastic distortion and a local core energy, due to lattice distortion that cannot be explained by linear elastic theory. Elastic force at a node is due to segments connected to the node, the imposed stress field (if one exists) and the field due to all other segments. Elastic forces weaken as  $1/r$  i.e., not very fast, and consequently, cannot limit their effect to a small neighborhood. This leads to global interaction between nodes and segments, which is mediated by a Fast Multipole Method (FMM). The use of FMM and the attendant recursive discretization of the domain into a dyadic mesh introduces (1) certain restrictions on how domain decomposition can be performed and (2) a set of numerical parameters (e.g. order of multipole expansions) that are specific to the FMM.

The force at a node is translated into a velocity using a mobility function / dislocation velocity submodel ( $\mathcal{M}_1$ ), obtained by fitting a “constitutive law” to molecular dynamics (MD) data. The velocity, in turn, is used to update nodal and segment positions using a backward-Euler time-integrator, with an adaptive time-step. The timestep is kept small enough such that a segment does not cross over to a new FMM subcell in the midst of the timestep.

## 4 Outline of the uncertainty quantification study

A categorization of ParaDiS inputs was discussed in Sec. 1. These inputs can assume a range of values, only some of which may be due to a lack of knowledge of their true values. Often, the range reflects the diversity of uses or operational conditions that a metal/material may be subjected to. However, the accuracy of the DD model (and its submodels) can vary significantly throughout these ranges (due to shortcomings of the model itself), leading to inaccuracies and uncertainty in the predictions. In this section, we outline the ranges over which these parameters can vary.

### 4.1 Uncertainties in Type I parameters

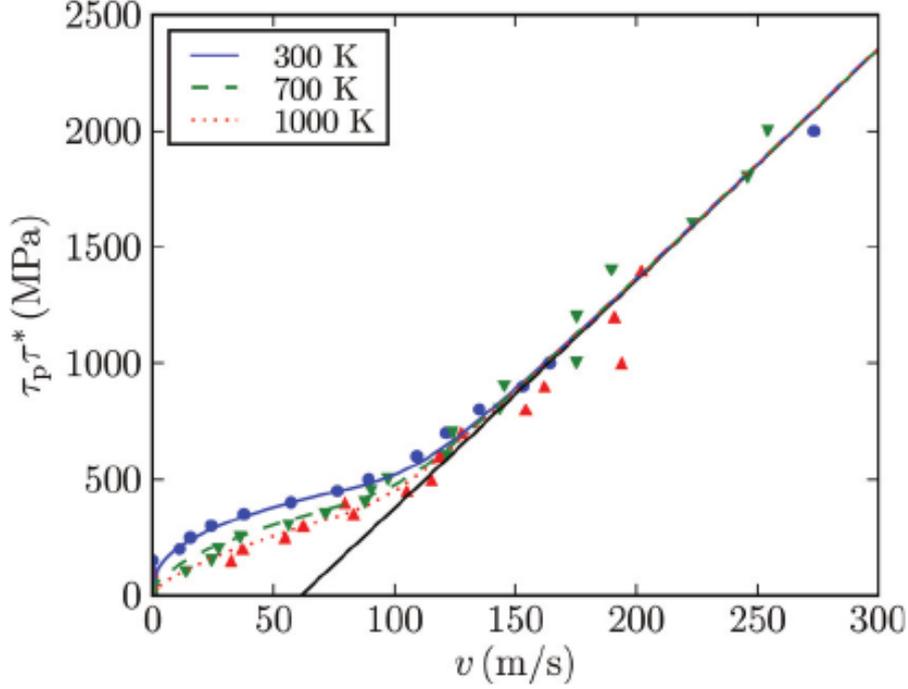
Type I parameters are parameters of a material's bulk properties. They are generally known with a high degree of precision. For  $T_a$ , Poisson ratio, Burger's magnitude and Young's modulus (`pois`, `burgMag`, `YoungModulus`) are known to 4 decimal places i.e., there is virtually no uncertainty in them and can be dropped from the UQ study. An exception could be made for Young's Modulus, since metals are anisotropic and one uses a volume-averaged value in ParaDiS .

The parameters `rc`, `MobClimb`, `MobLine` are numerical models that remove singularities in the existing model, preserve stability of the numerical scheme and allow linear elastic theory to be employed. Their values are fixed by trial and error, at a value that is as close to zero as possible, but which ensures stability of the model without impacting the model outputs. Again, these parameters can be dropped from the UQ study.

The parameters `MobGlide` and `MobPieirls` are obtained by fitting models to data from atomistic simulations. These parameters can be estimated to within a certain degree of accuracy, similar to Type III parameters. In Sec. 4.3 we describe how one estimates uncertainties in ParaDiS parameters (specifically, Type III parameters) that are obtained by fitting models to microscale simulation data. The same approach could be adopted for these two parameters.

### 4.2 Uncertainties in Type II parameters

Three inputs  $T$ ,  $\dot{\epsilon}$  and  $G$ , the shear modulus acting as a proxy for pressure, constitute Type II parameters. These are set by the operating conditions of the device made of Ta, and are thus entirely determined by the use-case. If there is interest in explaining the dependence of ParaDiS outputs on these parameters, they can be modeled as independent, uniform distributions. The lower and upper bounds for the  $(T, \dot{\epsilon}, P)$  parameters are  $(200, 1000)\text{K}$ ,  $(10^2, 10^5)\text{sec}^{-1}$  and  $(1, 30)\text{ GPa}$ .  $P$  is the pressure.



**Figure 1.** Fig. 3 from [1]. Blue and green symbols were read off with a graph paper and used as MD “observations” in the ParaDiS input characterization.

### 4.3 Uncertainties in Type III parameters

Type III parameters are subject to “true uncertainty”, i.e., we do not know their values with absolute certainty. These parameters are used by the dislocation velocity submodel ( $\mathcal{M}_1$ ; also called mobility function) used in ParaDiS. In this section, we characterize these Type III parameters (6 in number) with a joint distribution. These parameters appear in the “constitutive law” that describes the velocity of an isolated dislocation in a stress field, and is obtained via MD simulations. The model is obtained by fitting to MD  $\sigma_e - v$  data. The MD “observations” were read off by laying a graph paper on Fig. 3 in [1] (also, Fig. 1). The characterization is being done only for tantalum.

The model fitted was

$$\begin{aligned}
 \tau^* &= (\tau_T^q + \tau_D^q)^{1/q} \\
 \tau_T(v, \theta) &= \alpha_0 \exp\left(\frac{\theta}{\theta_T}\right) \left\{ \exp\left[\left(\beta_0 + \frac{\theta}{\beta_T}\right) \ln\left(\frac{v}{c_0} + \tilde{v}_0\right)\right] - \exp\left[\left(\beta_0 + \frac{\theta}{\beta_T}\right) \ln(\tilde{v}_0)\right] \right\} \\
 \tau_D(v, \theta) &= \max\left(\frac{(\chi_0 + \chi_{0,1}\theta + \chi_{0,2}\theta^2)(v/c_0 - \chi_1)}{\sqrt{1 - v^2/c_0^2}}, 0\right)
 \end{aligned} \tag{1}$$

**Table 2.** Mean, standard deviation and bounds of independent truncated normal prior distributions for the parameters  $\mathbf{p}$ .

Parameters	$\beta_0$	$\beta_T$	$\alpha_0$	$\alpha_T$	$\chi_0$	$\chi_1$
mean	0.25	2500	2.5	1500	50	0
std. dev	1.0	10,000	100	10000	30	1
nominal val	0.2757	2548	2.364	1127	37.166	0.03075
lower bound	0.05	500	0.1	500	1	-10
upper bound	0.5	5000	10	5000	100	10

The data consists of  $\tau^*\tau_p$ , where  $\tau_p$ , the Pierls stress at atmospheric pressure is tabulated in [1]. For tantalum,  $\chi_{0,1}$  and  $\chi_{0,2}$  are zero (not in the model at all),  $v_0 = 10^{-9}$  and  $q = 5$ . The net set of parameters are  $\alpha_0, \alpha_T, \beta_0, \beta_T, \chi_0$  and  $\chi_1$ . The parameters  $\chi_{0,1}, \chi_{0,2}$  are for the phonon-drag region, which is not a practically useful regime. A total of 34  $\sigma_e - v$  data points (10 for 300K and 14 for 700K) were used to estimate the 6 dimensional distribution.

If  $\mathbf{v}$  is the data vector, the inverse problem is formulated as

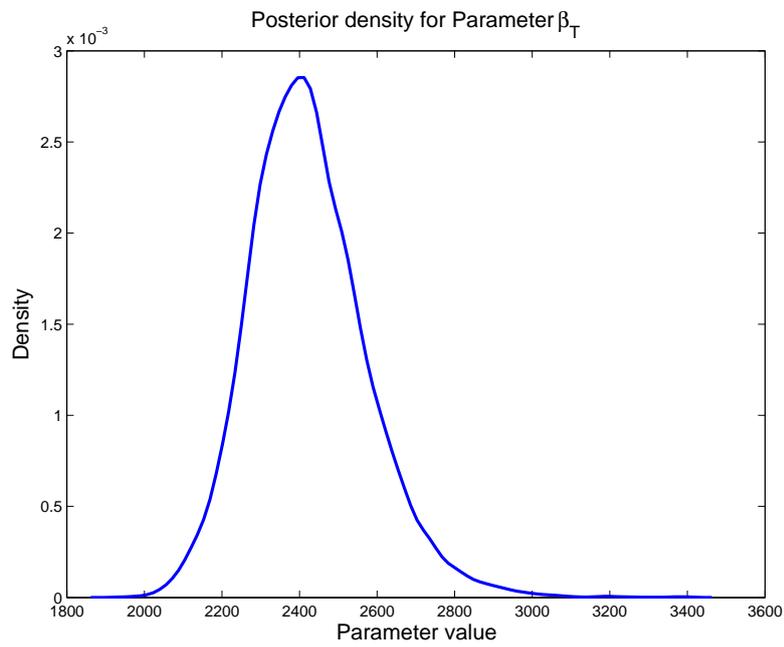
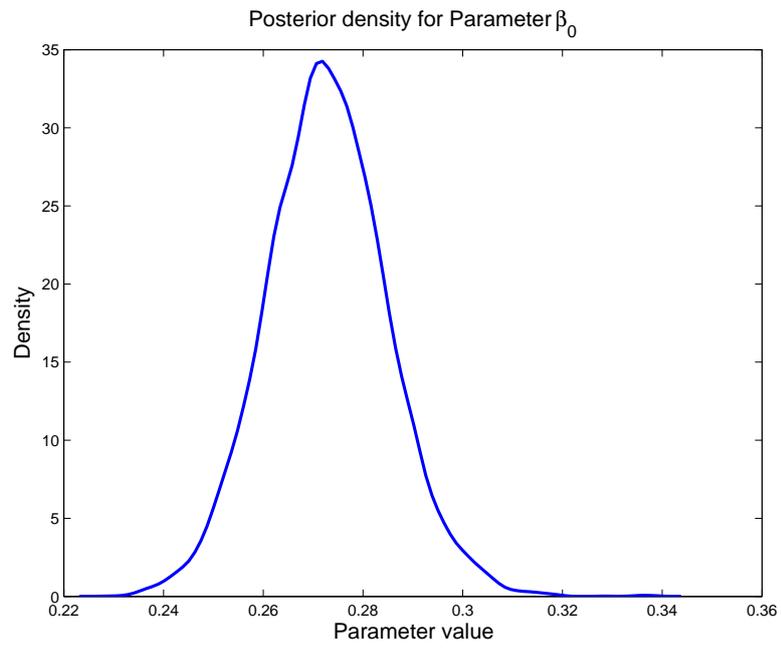
$$\mathbf{v} = \mathbf{f}(\boldsymbol{\sigma}_e; \mathbf{p}) + \boldsymbol{\eta}$$

where  $\mathbf{p} = \{\alpha_0, \alpha_T, \beta_0, \beta_T, \chi_0, \chi_1\}$ . By construction, the i.i.d. error  $\boldsymbol{\eta} \sim N(0, \boldsymbol{\sigma}_0^2)$  and  $\boldsymbol{\sigma}_0$  is also a parameter to be estimated. Note that the error  $\boldsymbol{\eta}$  is an amalgamation of the ‘‘jitter’’ in the MD simulations (variation from run to run at the same input parameters) and the model inadequacy. The parameters were estimated using DRAM [6]

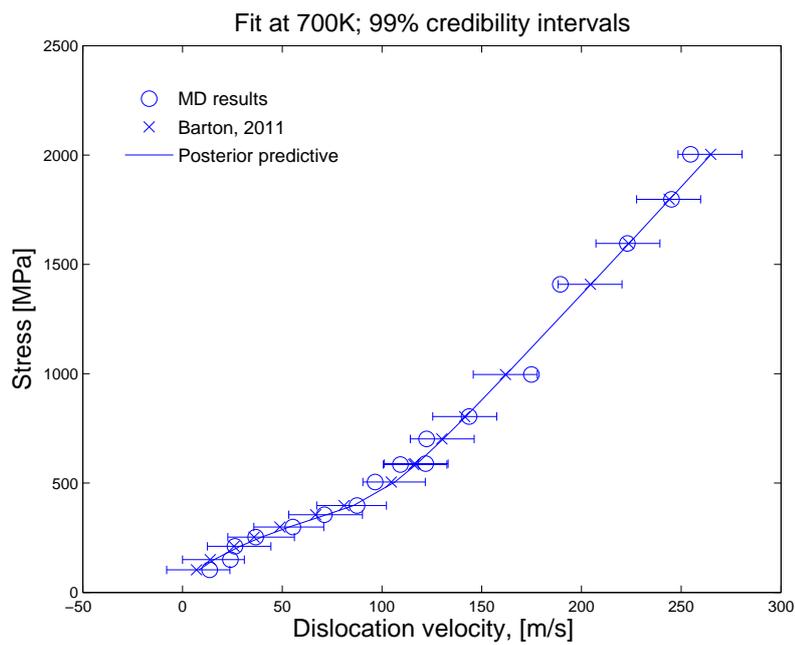
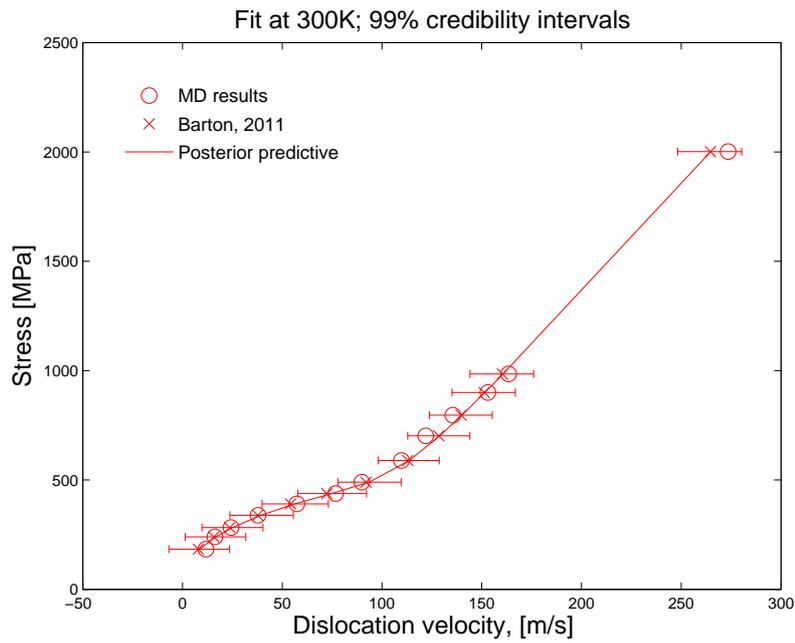
Independent truncated Gaussian priors are used for the 6 parameters. They are in Table 2. The standard deviation of  $\boldsymbol{\eta}$ ,  $\boldsymbol{\sigma}_0$ , which is also estimated, has an inverse chi-square distribution as a prior i.e.,  $\boldsymbol{\sigma}_0^{-2} \sim \chi^2(10, 4)$ .

### 4.3.1 Estimating $\beta_0$ and $\beta_T$

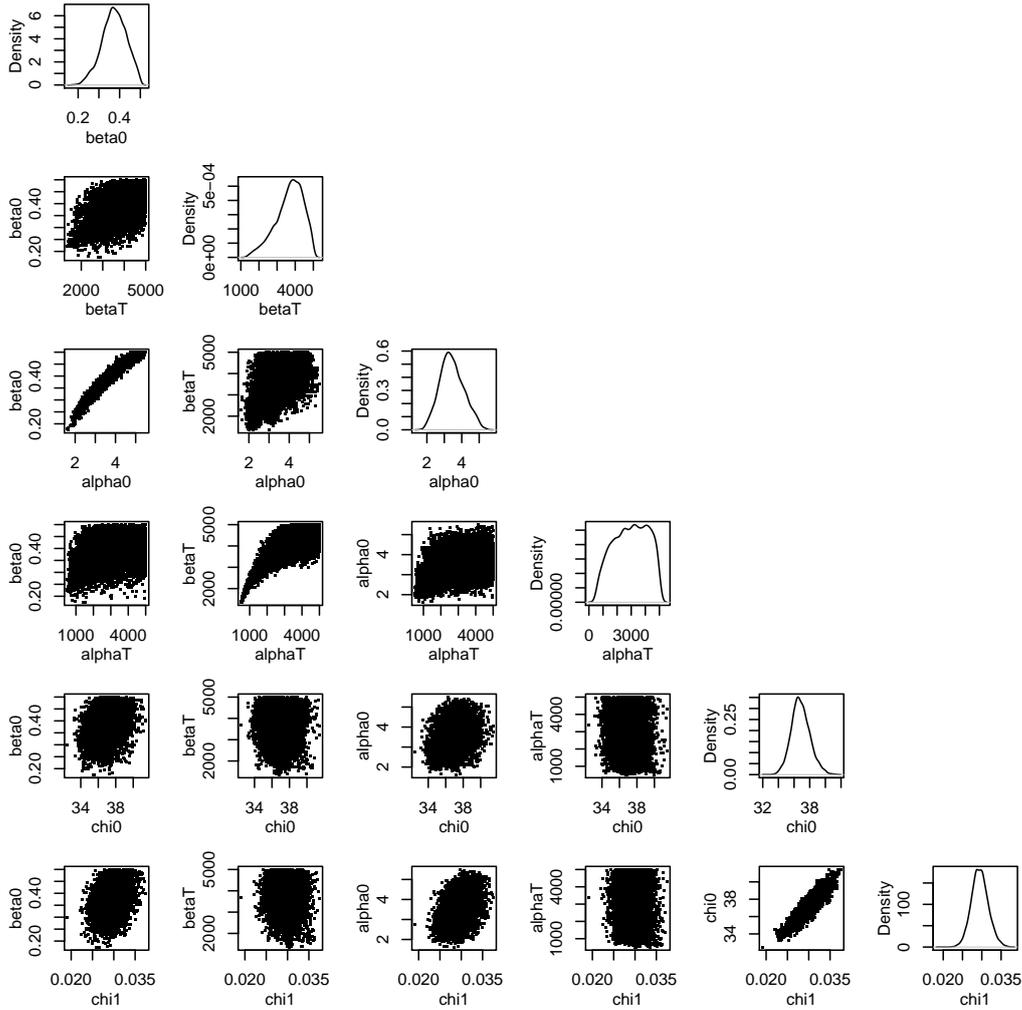
We conjecture that the most important parameters would be  $\beta_0$  and  $\beta_T$ , since they appear inside an exponential. Fig. 2 plots the 1D distribution of  $\beta_0$  and  $\beta_T$  (by marginalizing out the other parameter), when all other parameters are held constant at their nominal values (see [1]). In Fig 3 we plot the posterior predictive check. We see that the observations are contained inside the error bars, but the scatter in the posterior predictive test is greater than the data (most of the MD data lie near the median line). The posterior predictive check, therefore, fails.



**Figure 2.** Marginalized posterior distribution of  $\beta_0$  (top) and  $\beta_T$  (bottom).



**Figure 3.** Posterior predictive tests for 300K (top) and 700K (bottom), when 2 parameters are left uncertain. The tests do not match the variability in the data. The solid line is the median from the posterior predictive tests, and the error bars show the 99% credibility interval.



**Figure 4.** Pairwise plot of the 6 parameters being estimated.

### 4.3.2 Estimating six uncertain parameters

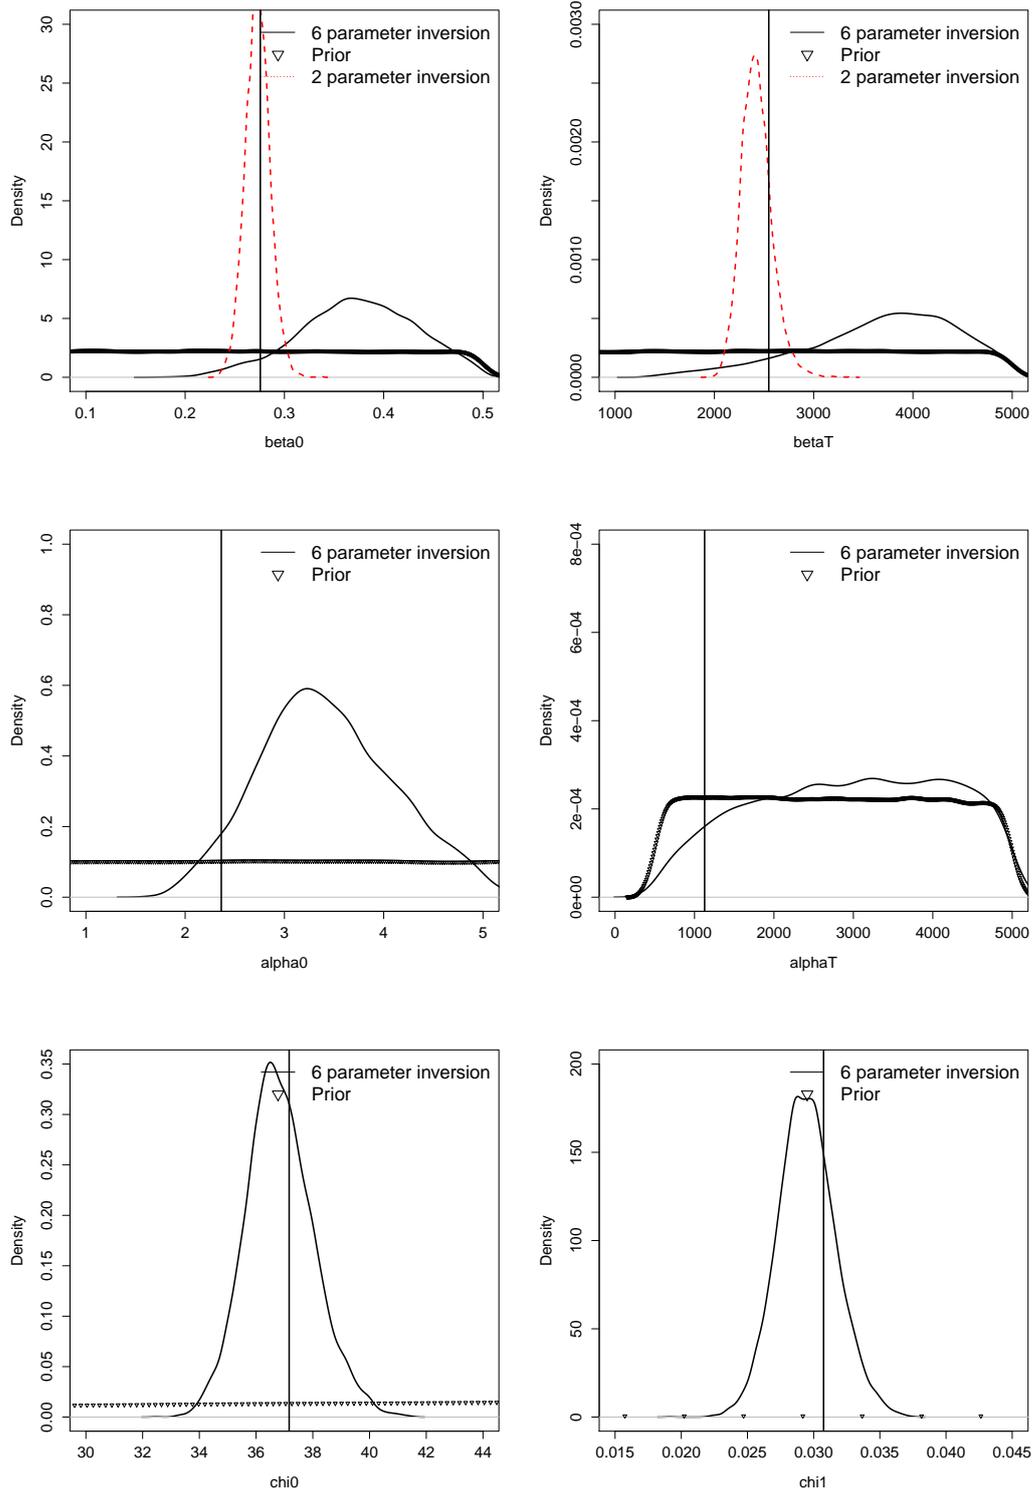
We now let all 6 parameters be uncertain and attempt to reproduce the variability in the MD data. In Fig. 4, we plot marginalized posteriors from the fit. We see significant correlations between  $\beta_0 - \alpha_0$ ,  $\beta_T - \alpha_T$  and  $\chi_0 - \chi_1$ . In Fig. 5, we plot the prior and posterior distributions and in case of  $\beta_0$  and  $\beta_T$ , comparison with the 2-parameter inversion. We see that the posterior distribution for  $\beta_0$  and  $\beta_T$  are very different when performed with 2 versus 6 uncertain parameters. This is because of their correlation with  $\alpha_0$  and  $\alpha_T$ . The PDF of  $\alpha_T$  shows very little specificity; due to its correlation with  $\beta_T$ , the PDF of  $\beta_T$  also broadens beyond the PDF observed for the 2-parameter case. The PDF for  $\alpha_0$  is quite broad too; consequently, the PDF of  $\beta_0$  broadens beyond the 2-parameter study, though not as much as  $\beta_T$ . The nominal values of  $\beta_0$  and  $\beta_T$  are very close to the MAP values obtained from the 2-parameter study.

In Fig. 6 we look at posterior predictive tests generated using the joint distribution depicted in Fig. 4. Again the posterior predictive checks bear out. Further, due to the larger dimensionality, the 6-parameter model fits better.

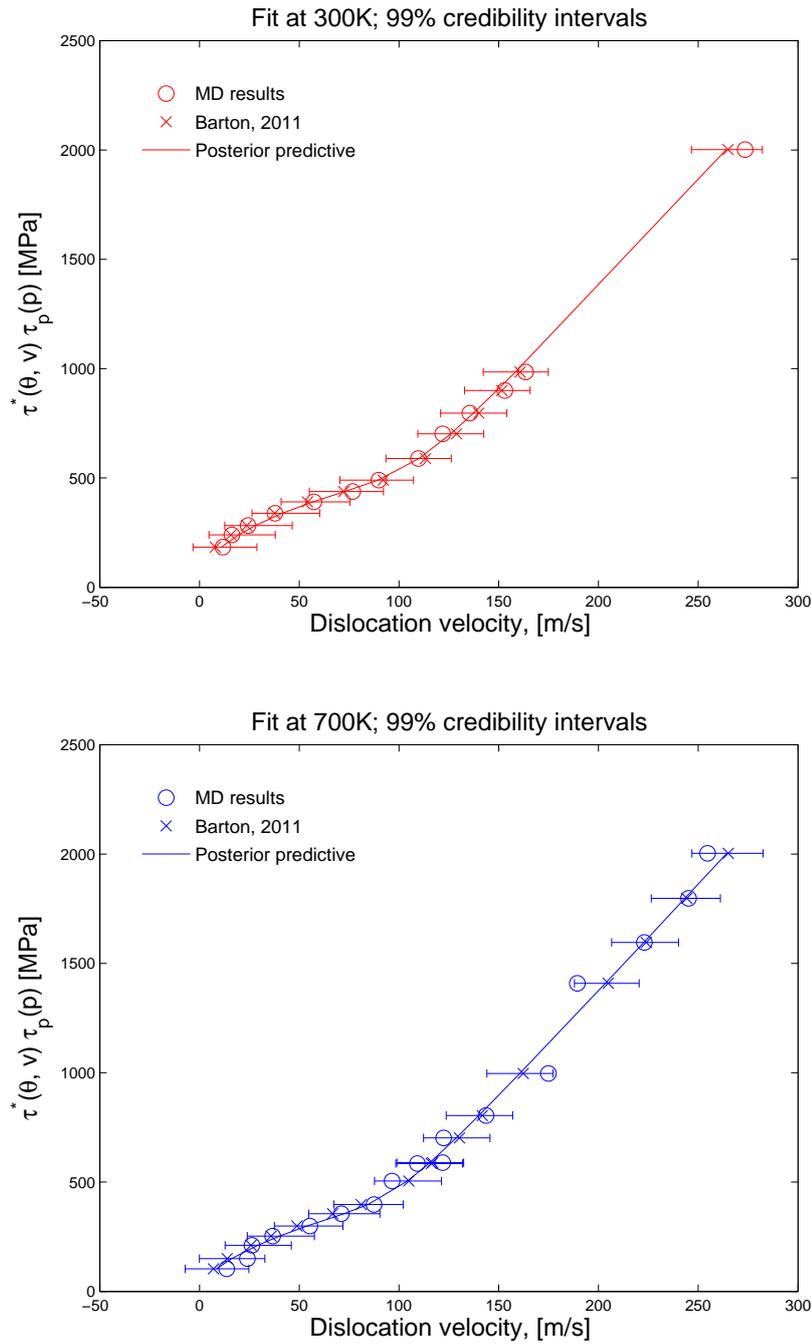
In Fig. 7, we plot the model error  $\eta \sim N(0, \sigma_o^2)$ . There is not much difference whether 2 or 6 uncertain parameters are used.

Finally, we address the quality of calibration of the 2-parameter and 6-parameter models using verification rank histograms [9]. We combine the predictions of the dislocation velocity in Fig. 3 and the observations, sort them and determine the rank the observations in the sorted set. Since there are 1000 predictions, the ranks of 30 observations should vary evenly between 1 and 1001, if the model is well calibrated (concentration of the ranks around the middle of the range indicates an under-dispersed calibration; concentration at the extremities an over-dispersed calibration). We plot the histogram of ranks, binned into 30 equally-sized bins, in Fig 8, above, for the two parameter model. Its counterpart, developed using the 6-parameter model, is plotted in Fig. 8 (below). Since we have 30 observations and 30 equally-spaced bins, an ideal calibration would result in all bins with one observation each. We clearly see that this is not achieved for the 2-parameter model (Fig. 8, top) but is more common for the 6-parameter model. While this indicates that the 6-parameter model is better calibrated, it does not automatically follow that it is more predictive.

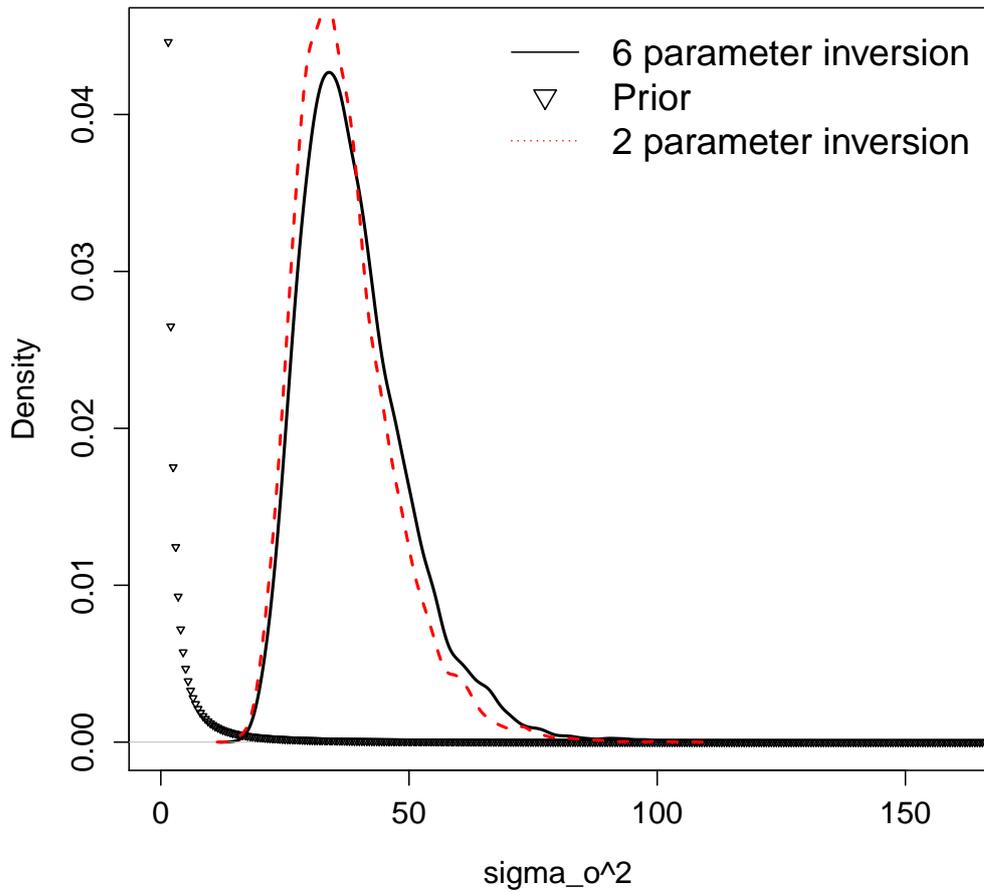
To summarize, the 6 parameters of Eq. 1 can be estimated as distributions from MD data. Our conjecture that it may be sufficient to calibrate just two parameters ( $\beta_0, \beta_T$ ) was shown to be wrong; Fig. 4 shows significant correlation between these two parameters and others in Eq. 1. The verification rank histograms show that the 6-parameter model is better calibrated. However, the degree of mismatch  $\sigma_o$  was computed to be about the same in both the two- and six-parameter model calibrations, indicating that adding an extra four calibration parameters did not change the predictive skill of the model significantly.



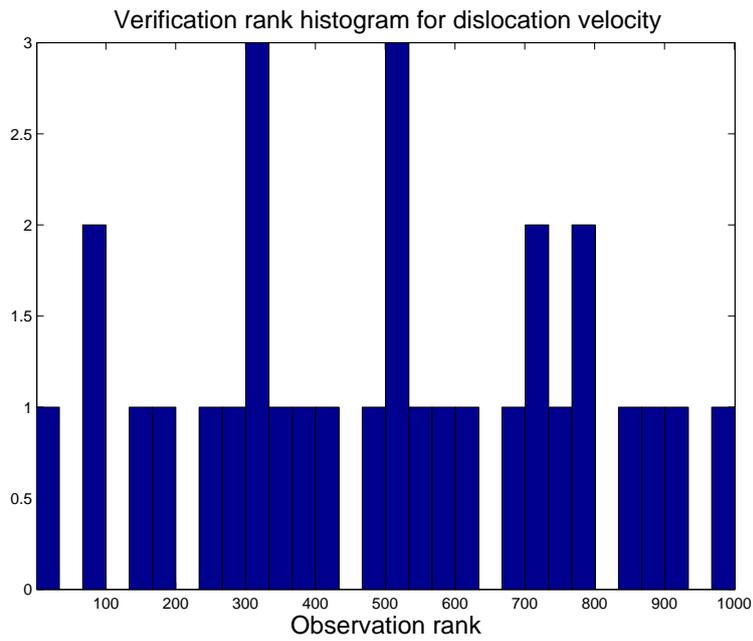
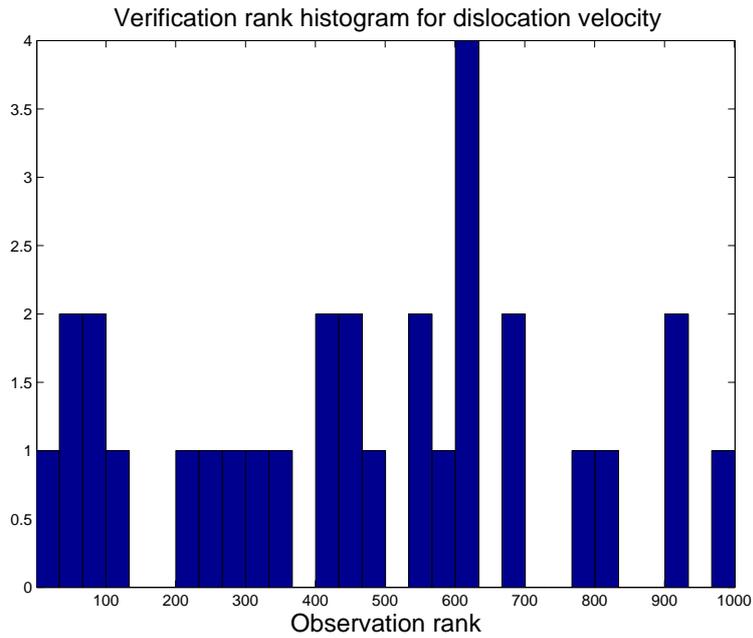
**Figure 5.** Prior and posterior distributions for  $\beta_0$ ,  $\beta_T$ , (top)  $\alpha_0$ ,  $\alpha_1$ , (middle)  $\chi_0$  and  $\chi_1$  (bottom). The nominal values, currently used in ParaDiS, are plotted as vertical lines. Red lines denote the distributions computed with just  $(\beta_0, \beta_T)$  assumed uncertain (Sec. 4.3.1).



**Figure 6.** Posterior predictive tests for 300K (top) and 700K (bottom). We see that we capture the variability of the MD data. The solid line plots the median of the posterior predictive test and the error bars cover the 99% credibility interval.



**Figure 7.** Posterior for  $\sigma$ , using 6 uncertain parameters.



**Figure 8.** Verification rank histograms for the 2-parameter model (top) and the 6-parameter model (bottom). Thirty observations are binned into 30 equally-spaced bins. We clearly see that the calibration of the 2-parameter model is under-dispersive, whereas the 6-parameter model displays a good calibration.

## 4.4 Designing a computationally inexpensive ParaDiS problem

The need to run ParaDiS for a number of parameter settings requires one to construct a problem configuration that is small enough to be run repeatedly. This can be achieved with a small domain, a coarse mesh and a shorter duration over which DD is simulated (alternatively, a smaller final strain). We can also relax some of the numerical parameters, at the cost of incurring some numerical errors. This necessarily means forgoing some accuracy in the simulated physics. Since the data so generated will be fitted with  $\mathcal{M}_2$ , we expect that some of these inaccuracies will be reflected in the fit. However, the bulk of the uncertainty in  $\mathcal{M}_2$  predictions are expected to be due to physical parameters (Type I, II & III) rather than the numerical ones (Type IV) which govern how quickly the simulation can be run.

We have constructed such a configuration for Ta; further, it was used to generate the ParaDiS results in Sec. 4.6. The domain is a cube,  $2.5\mu$  on each side. The domain is divided into a mesh with 8 cells on each side; these are subsequently refined, if needed. The time-integrator is a second-order backward-Euler method. Timesteps were nominally set to  $10^{-14}$  seconds, but allowed to adapt to as much as  $10^{-6}$  seconds, if needed. The multipole expansion is performed only up to second order. Segment lengths are allowed to vary between 50 and 200 nm. The problem runs on about 64 cores for 3 weeks (including time spent in queues etc).

However, while we have constructed such a small configuration, a thorough analysis of the various numerical parameters on the fidelity of the ParaDiS outputs has not been performed. Prior to using the small configuration in uncertainty quantification studies, we will evaluate the impact of Type IV (i.e., numerical) parameters on ParaDiS predictions and ensure that the uncertainties introduced by them are substantially smaller than those due to physical parameters of ParaDiS.

## 4.5 Sampling strategy for ParaDiS

Given the discussion in Sec. 2.3, it is clear that a drastic reduction in the dimensionality of the UQ problem will be required. The dimensionality reduction is also dependent on the computational resources one has access to. However, we assume about 100–500 ParaDiS runs may be possible. This, in turn, implies that we may consider 5-10 uncertain ParaDiS input parameters.

We do not expect any reduction in the Type II parameters which span operational conditions.

Some Type I parameters will be discarded in consultations with experts. Crystallographic parameters do not vary significantly, whereas those dependent on grain size/distribution may be affected by manufacturing and metallurgical processes and display a larger degree of uncertainty. Type I parameters hold the potential for the largest reduction in dimensionality,

but the process may be somewhat ad-hoc.

Type III parameters may be discarded based on their ability to affect the predictions of the mobility function  $\mathcal{M}_1$ . The discussion in Sec. 4.3 shows that the prediction error  $\sigma_0$  obtained with two uncertain parameters  $(\beta_0, \beta_T)$  is virtually the same as when all 6 parameters were varied, where significant correlations between parameters were also observed. Performing the same fitting with different pairs of parameters would reveal whether  $(\beta_0, \beta_T)$  were indeed unique, since a 2-parameter model had the same predictive skill as a 6-parameter one. If so, only  $(\beta_0, \beta_T)$  could be retained as uncertain parameters, with the rest (the non-sensitive parameters) set to their nominal values.

An alternative method, which would not allow more dimensions, but may allow simpler sampling, is random sampling e.g., Latin Hypercube or quasi-Monte Carlo. In such a case, surrogate models would have to be fitted via least-squares regression to ParaDiS output. The limits on the complexity of the polynomial form being fitted, as discussed in Sec. 2.3, would also hold. However, a random distribution of sample points would hold other advantages. While a PC (or any polynomial) basis set would provide us with a “trend” model, one could also fashion a “correction” to the PC/polynomial model predictions by interpolating from nearby sample points using, for example, kriging or inverse distance weighting. However, this option can only be evaluated once a few ( $O(10)$ ) ParaDiS runs have been made and one has an approximation of the response of ParaDiS to parameter perturbations.

## 4.6 Estimating model uncertainty in $\mathcal{M}_2$

The output from a ParaDiS run is fitted with a model (Eq. 11 in [1]; also called  $\mathcal{M}_2$ ) which is reproduced here,

$$\dot{\rho} = R\dot{\epsilon} \left( 1 - \frac{\rho}{\rho_{sat}(\dot{\epsilon})} \right), \quad (2)$$

where  $\dot{\epsilon}$  is the strain rate,  $R$  is a constant,  $\rho$  is the dislocation density and  $\rho_{sat}(\dot{\epsilon})$  is its saturated value (the asymptotic value reached at large strain). The saturated value depends on the strain rate and the temperature. This simple model is fitted to ParaDiS output only after the material has started undergoing plastic deformation (which is determined from the stress-strain curve, also a ParaDiS output).

The fit of Eq. 2 employs “data” that is in fact the output of a deterministic computational model, namely ParaDiS. The difference between ParaDiS outputs and  $\mathcal{M}_2$  (Eq. 2) predictions cannot be ascribed to noise (there is no noise in ParaDiS outputs). It is feasible, as is done in the literature, to model this difference nonetheless with a statistical model. However, there is no clear justification of any specific form of this model, as there is no data-noise that could be employed to inform this construction. This, therefore, introduces additional modeling artifacts. Moreover, when the Bayesian fitting is done in this manner, the resulting uncertainty in the fitted model parameters, when pushed forward through this model, typically leads to a pushed-forward posterior density on model outputs that is narrower than the range of discrepancy with the ParaDiS results. This is due to the noise

model employed in the fitting, which typically ends up being calibrated to explain a good portion of this discrepancy, but yet is not part of the upscaled model. A posterior predictive, resulting from the fitted model combined with the noise model, would in fact be expected to be consistent with the spread of the data, but not the fitted model itself. Accordingly, expressing here the desire that the pushed-forward posterior on fitted model outputs be consistent with the data spread, and hence with the discrepancy from ParaDiS outputs, we propose the following alternate approach.

We propose to model  $(R, \rho_{sat})$  as random variables with an unknown, but parameterized, joint probability density distribution e.g., one could model the random variable behind this distribution using polynomial chaos (PC) expansion. The use of a distribution to describe  $(R, \rho_{sat})$  results in a probabilistic prediction of  $\rho(t)$  using Eq. 2 i.e., the prediction itself is a distribution. We would like the  $(R, \rho_{sat})$  distribution to be calibrated such that the  $\rho(t)$  distribution fits, and is consistent with the spread of, the data. This requirement forms the basis for the following model fitting procedure.

To perform a preliminary analysis, we model the distribution of  $(R, \rho_{sat})$  as a bivariate Gaussian i.e.  $(R, \rho_{sat}) \sim MVN(\mu, \Gamma)$ , where  $\mu$  is a vector of size 2 and  $\Gamma$  is a  $2 \times 2$  covariance matrix. There are 5 unknowns – the 2 members of  $\mu$  and the 3 constituents of  $\Gamma$ . Note that since  $\Gamma$  is symmetric positive definite,  $\Gamma_{12} = \Gamma_{21}$ . The 5 unknowns are estimated from ParaDiS outputs. For a proposed  $(\mu, \Gamma)$ , we can obtain  $N$  samples of  $(R, \rho_{sat})$  pairs from the distribution and trace the evolution  $\rho_i(t), i = 1 \dots N$ , for each sample  $i$ . At any time  $t$ , the model predictions form a distribution with a mean  $\bar{\rho}(t)$  and a standard deviation  $\sigma_\rho(t)$ .

Let  $\rho^{ParaDiS}(t_j)$  be the ParaDiS prediction of dislocation density at time instant  $t_j$ . Let  $\bar{\rho}(t_j; \mu, \Gamma)$  be the mean of the predictions (at time  $t_j$ ) using Eq. 2, corresponding to a distribution of  $(R, \rho_{sat})$  obtained from a  $MVN(\mu, \Gamma)$ . The values of  $(\mu, \Gamma)$  are estimated by the following optimization

$$\begin{aligned} & \underset{(\mu, \Gamma)}{\text{minimize}} && \xi^2(\mu, \Gamma) \\ & \text{subject to} && (|\xi(\mu, \Gamma)| - \sigma_\rho)^2 \leq \varepsilon_s \end{aligned} \quad (3)$$

where  $\xi^2 = \sum_j (\rho^{ParaDiS}(t_j) - \bar{\rho}(t_j; \mu, \Gamma))^2$  and  $\varepsilon_s$  is a small number. The optimal values of  $(\mu, \Gamma)$  so obtained can be used to test the goodness of fit of the model, by generating predictions using Eq. 2 and comparing with ParaDiS output. We demonstrate this below.

We use a ParaDiS simulation corresponding to the conditions described in Sec. 4.4. The material is Ta, subjected to a strain-rate  $\dot{\varepsilon} = 10^5 \text{sec}^{-1}$ . Log-transformed  $\rho$ , rather than  $\rho$  itself is used in fitting Eq. 2. In Fig. 9 (top), we plot the evolution  $\rho^{ParaDiS}(t)$  (in black) as well as the stress-strain curve (in blue). The stress-strain curve is used to demarcate the plastic region (plotted in red). We fit Eq. 2 as a deterministic model to obtain point estimates of  $(R, \rho_{sat})$ ; its prediction is plotted in green. We see that Eq. 2 provides a good fit, but the ParaDiS output “wiggles” around it. The optimal point estimates are  $R = 9.57 \times 10^{15} \text{m}^{-2}$  and  $\rho_{sat} = 1.54 \times 10^{14} \text{m}^{-2}$ .

Eq. 3 is then used to estimate the 5 variables that constitute  $(\mu, \Gamma)$ .  $\rho_{sat}$  is bounded below at  $\rho(t = 0)$ , the initial condition. Thereafter, we sample  $(R, \rho_{sat})$  pairs from the resulting

distribution and plot them as a scatterplot in Fig. 9 (bottom). The fitting procedure reveals the existence of multiple local minima, which, nevertheless are reached at the same value of  $\Gamma_{22}$ . The log-likelihood values of these minima are all approximately the same. Thus, we find that only  $\Gamma_{22}$  can be estimated with any degree of confidence; the rest of elements of  $\Gamma$  are uncertain and also do not make much of a difference to predictions of  $\rho$  using Eq. 2.

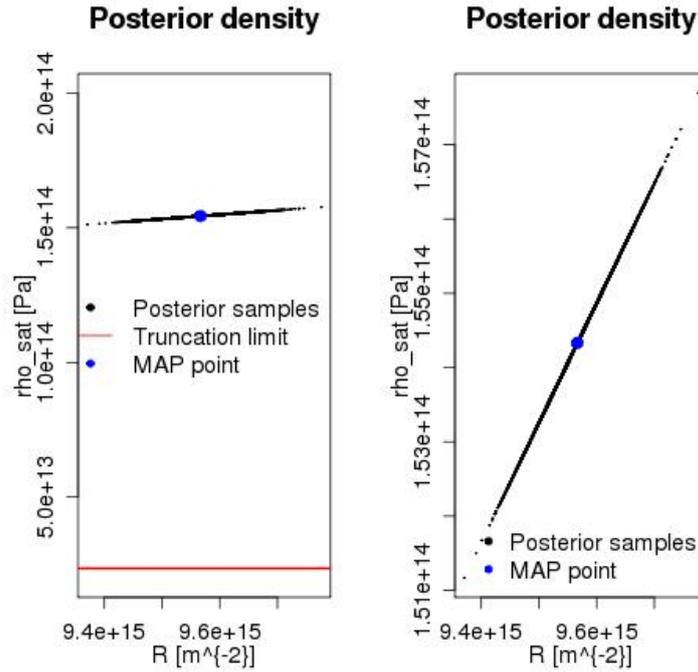
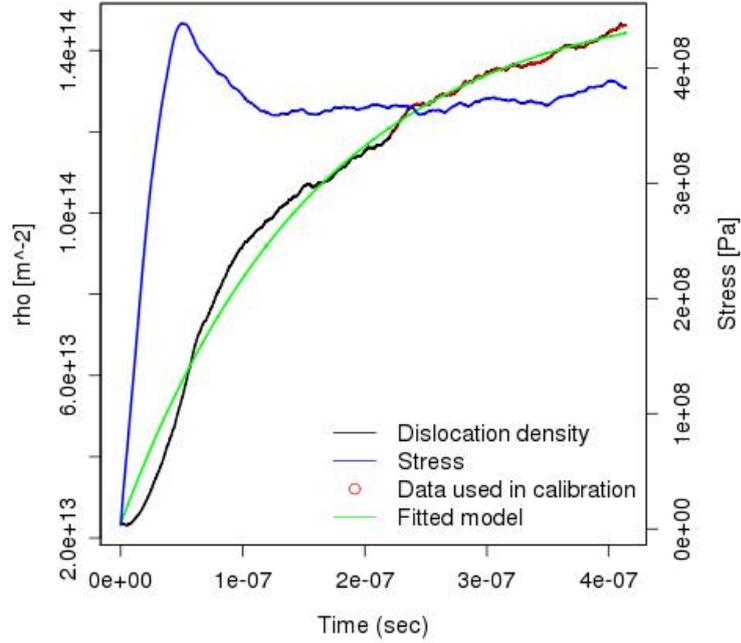
In Fig. 10, we use the fitted distribution  $MVN(\mu, \Gamma)$  to generate 100 predictions of  $\rho(t)$ . These are plotted with green symbols in Fig. 10 (top). We see that the “fan” of predictions using Eq. 2, with  $(R, \rho_{sat})$  modeled as a bivariate Gaussian, spans the ParaDiS output, plotted in red. In Fig. 10 (bottom) we plot the distribution of predictions  $\rho(t)$  at 6 different points in time (each with a different color). The dashed lines are the ParaDiS results. We see that there is no clear trend - the probabilistic predictions are neither consistently below nor above the ParaDiS values. Rather, we see a correlated behavior, where the fitted model first overpredicts (early time and lower  $\rho$  values), then draws close to the ParaDiS values (blue and red curves), followed by more over-prediction (green and magenta curves) and finally an under-prediction (bold magenta). This meandering of the curve about the middle of the distribution is also seen in Fig. 10 (top).

To summarize, one can model  $(R, \rho_{sat})$  as a distribution and use Eq. 2 for predicting  $\rho(t)$ . We have demonstrated a method to estimate this distribution, by modeling it as a bivariate Gaussian; other models can also be adopted and fitted using the same procedure. We see that the distribution is narrow and is sufficient to “cover” the shortcomings of Eq. 2 when fitted to ParaDiS data, inasmuch that the distribution of Eq. 2 predictions contain the data used for fitting. Note that this distribution is derived from *one* setting of ParaDiS inputs i.e., it does not reflect any input uncertainty on the part of ParaDiS. The uncertainty captured by the  $(R, \rho_{sat})$  distribution is purely model uncertainty i.e., it reflects the shortcomings of  $\mathcal{M}_2$  (Eq. 2) in representing ParaDiS outputs.

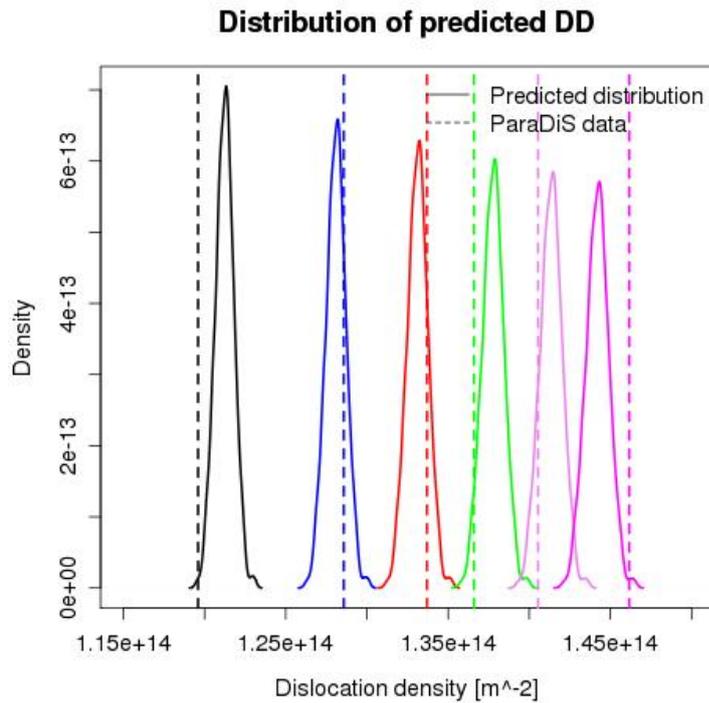
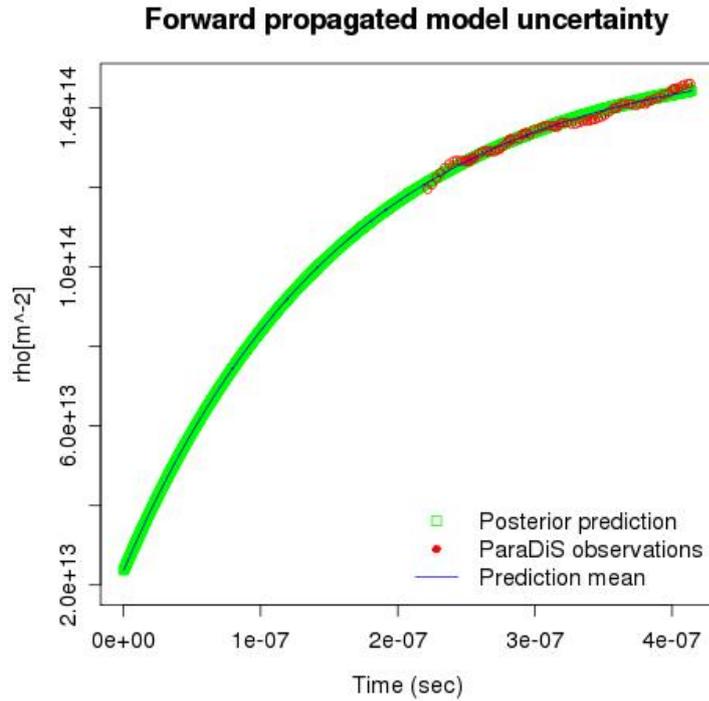
## 4.7 Execution plan and deliverables

The execution of a ParaDiS UQ study would consist of the following steps:

1. *Dimensionality reduction:* The first step of the study would target dimensionality reduction. Insensitive Type I parameters would be identified in consultation with experts. Sensitivity analysis of the dislocation mobility function, (after being fitted to MD data) will be used to reduce the dimensionality of  $\mathcal{M}_1$ . We will aim for 5–10 uncertain parameters.
2. *Sampling design:* The ParaDiS parameters will consist of a mixture of independent parameters (Type I and II) and correlated ones (Type III). While this poses no problems in case of random sampling, designing a CC grid with correlated parameters is expected to be a challenge and require significant resources. This may also be an argument against adopting sparse quadratures for our sampling study.



**Figure 9.** Data for fitting Eq. 2 and the resulting fit. Top: ParaDiS output in black (plotted against the left Y-axis) and the stress-strain curve (blue; plotted against the right Y axis), showing the plastic region. The red region is the plastic-deformation data used for fitting Eq. 2. The green line is the fitted model, using optimal values of  $R, \rho_{sat}$ . Bottom left:  $(R, \rho_{sat})$  samples from the fitted distribution. We see that there is very little uncertainty in the fit of  $\mathcal{M}_2$ ; the distribution is very narrow, but  $(R, \rho_{sat})$  clearly assume correlated values. Bottom right: A zoom of the fitted (very narrow) Gaussian.



**Figure 10.** Predictions with the fitted model. Top: Using values of  $(R, \rho_{sat})$  sampled from the distribution plotted in Fig. 9 (bottom), we prediction an ensemble of 100  $\rho$  evolutions, plotted in green. The ParaDiS output is plotted in black. Bottom: Distribution of predictions  $\rho(t)$  (arising from Eq. 2 and  $(R, \rho_{sat}) \sim \text{MVN}(\mu, \Gamma)$ ) at 6 different points in time (each corresponding to a different color). The dashed lines are the ParaDiS results. The black plot is at the earliest time, the bold magenta at the latest time.

3. *Execution of the sampling design:* This is expected to be the most time-consuming part of the study, as it involves nursing  $O(100)$  ParaDiS executions to conclusion and fitting  $\mathcal{M}_2$  to their individual outputs.
4. *Construction of the surrogate model:* The steps described above will result in distributions of  $(R, \rho_{sat})$ , with each distribution corresponding to a ParaDiS run. These distributions may be represented by a  $(\mu, \Gamma)$  set, if a multivariate Gaussian distribution is found to be adequate, or a set of PC coefficients, if we take recourse to PC-based modeling of the joint  $(R, \rho_{sat})$  distribution. Surrogate modeling will take the form of establishing polynomial maps between the 5–10 unknown parameters and the PC coefficients (or  $(\mu, \Gamma)$ ). Depending upon the complexity of the resulting surface, interpolation (as described in Sec. 2.3) may also be required.
5. *Global sensitivity analysis and UQ studies:* The final step of the study would be to use the surrogate to establish (1) the most sensitive parameters, (2) the uncertain outputs, at specified Type II (“operational”) parameters values, by marginalizing over the uncertain parameters. This would yield a distribution over  $(R, \rho_{sat})$ , which are input parameters for  $\mathcal{M}_2$ , a submodel for ALE3D. Its use in an ALE3D sensitivity study is, thereafter, conceptually simple.

A UQ study of ParaDiS would yield two tangible deliverables:

1. *An approach for performing UQ studies on complex models:* An UQ study on ParaDiS would design and demonstrate how UQ investigations could be performed on complex and computationally expensive models with numerous inputs. The method would be general, treating a mixture of independent and correlated model inputs (e.g., Type III parameters) and addressing model predictions which were probabilistic i.e., model outputs are probability densities. A key deliverable of this approach would be the sampling design necessary to cover the 10-dimensional input space of ParaDiS.
2. *A surrogate model for the upscaling toolchain:* The sampling study would establish a map between the input parameters of ParaDiS and the distribution of  $(R, \rho_{sat})$  e.g., if we were to retain our multivariate Gaussian model for the distribution of  $(R, \rho_{sat})$ , the mapping would target the 5 constituents of  $(\mu, \Gamma)$ . A surrogate model could relate the sensitive inputs to the 5 outputs; alternatively, one could marginalize out the Type I and Type III parameters to obtain a relationship between Type II parameters (“the operational conditions”) and the resulting distribution of  $(R, \rho_{sat})$ .

A surrogate model would have many practical uses. Primarily, one would use it to perform a sensitivity analysis, to establish the input parameters that most impact outputs; this would subsequently guide experiments or MD simulations that would be instrumental in reducing parametric uncertainty.

The approach, on the other hand, could also be used to create competing upscaled surrogate models, predicated on competing mobility functions  $\mathcal{M}_1$ . The surrogates could then

be used to perform model selection i.e., pick the mobility law most supported by evidence by comparing with data obtained from anvil or explosive tests. A computationally inexpensive surrogate would allow the model selection to be performed rigorously i.e., by means of Bayes factors [7] or by comparing the quality of predictions versus data [8, 9]. Both these methods require multiple invocations of the surrogate model, and hence the need for computational celerity.

## 5 Computational issues and resources

We conclude with a discussion of the computational resources needed for this UQ study.

- *Estimation of the impact of Type IV parameters:* The impact of Type IV (numerical) parameters of ParaDiS will be gauged, with a view of ensuring that their impact is smaller than the other uncertain parameters. Since there are 12 such parameters,  $O(10)$  runs will be required to accomplish this task. We will simply run ParaDiS (the configuration described in Sec. 4.4) at parameter setting half their nominal value to estimate the impact.
- *Executing the sampling study:* In this case, we expect that the computational resources will determine the number of ParaDiS runs that can be executed. We estimate that we can execute 5 runs a month; each run will execute on 64 processors. Thus over 2 years (the recommended duration of the UQ study), we expect to obtain about 120 ParaDiS runs. This is equivalent to 4–6 million CPU hours.
- *Type of computers required:* There are no specific requirements on the type of computers required to execute the UQ study. If access to a large number of cores is feasible, they will be used to simply execute more ParaDiS sample runs concurrently.
- *Team structure:* Executing the UQ study will require a tight collaboration with the LLNL ParaDiS team. Since the UQ study will target Ta as the test material, it will have to be performed with the LLNL version of ParaDiS, which will also have to be ported to the various machines where the numerical experiments will be performed.

This page intentionally left blank

## 6 Summary

This report outlines an UQ study of ParaDiS , a dislocation dynamics simulator, being developed in LLNL. ParaDiS simulates the evolution of dislocations (and related physical outcomes, like plastic strains and strain-hardening) in a crystal when subjected to a stress or strain field. It uses submodels, the most important of which is the dislocation mobility function,  $\mathcal{M}_1$  . It is derived from molecular dynamic simulations and contains uncertain parameters which are ParaDiS inputs. ParaDiS predicts the time evolution of dislocation density, which is fitted with an upscaling model  $\mathcal{M}_2$  , and used in continuum simulators like ALE3D. ParaDiS has inputs that are classified as (1) material parameters, (2) operational parameters (e.g., temperature, pressure etc.), (3) mobility function parameters and (4) numerical parameters. There are 20 physical parameters and 12 numerical ones. ParaDiS is expensive and fully resolved runs take  $O(10^7)$  CPU-hours.

Tantalum will serve as the test metal for the purposes of the study.

The UQ study seeks to establish the uncertainties in  $\mathcal{M}_2$  predictions due to uncertainties in ParaDiS input. Due to the high-dimensionality of the input space and the huge computational expense of ParaDiS , it is clear that (1) the study has to be performed using a less expensive configuration and (2) the UQ study can only be done for a handful of parameters. In Sec. 4.4, we describe a computationally inexpensive problem configuration that can be completed in 3 weeks of runtime on 64 cores. In Sec. 4.5 we discuss how the 20-dimensional input parameter space may be reduced to 5–10 for the UQ study and identify some of the parameters that need to be retained. The number of parameters to be retained in the UQ study was arrived at by our estimation that one could execute 100–200 ParaDiS runs over a period of 2 year; these runs would be sufficient to resolve the dependence of ParaDiS outputs only if the number of independent inputs were severely restricted (see discussion in Sec. 2.3)

The set of ParaDiS outputs will be fitted with  $\mathcal{M}_2$  . In Sec. 4.6, we show that  $\mathcal{M}_2$  is incapable of reproducing ParaDiS outputs unless we assume its parameters to be random variables. In the same section, we demonstrate a method, based on multivariate Gaussian modeling of their distribution, by which the random variables may be estimated from ParaDiS data. We also outline a second one, based on polynomial chaos modeling of  $\mathcal{M}_2$  's parameters, in case the first proves inadequate.

We expect that the study will require 4–6 million CPU-hours, on modest to large platforms (due to the sheer number of runs that we require). We recommend that the team performing the UQ study be tightly integrated with LLNL. Since Ta is the metal of interest, the study will require the use of the LLNL (rather than the public domain version) of ParaDiS . Further, significant porting effort will be required.

The UQ study will deliver an approach for performing UQ studies on computationally expensive models, with ParaDiS as an example. The approach can be also be used to discriminate between competing version of ParaDiS e.g., if they use different submodels.

It will also deliver a surrogate for ParaDiS which can then be used for validation studies (against anvil and explosive test data) or for performing sensitivity analysis under different operational (temperature, pressure or strain-rate) conditions.

## References

- [1] N. R. Barton et al. A multiscale strength model for extreme loading conditions. *Journal of Applied Physics*, 109(073501), 2011.
- [2] ParaDiS Homepage. <http://paradis.stanford.edu/>.
- [3] ALE3D Homepage. <https://wci.llnl.gov/codes/ale3d/>.
- [4] J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*. Springer, New York, 2004.
- [5] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [6] H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM-Efficient adaptive MCMC. *Statistics and Computing*, 16(4):339–354, 2006.
- [7] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman & Hall / CRC, Boca Raton, FL 33431, 2004.
- [8] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [9] Tilmann Gneiting, Fadoua Balabdaoui, and Adrian E. Raftery. Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):243–268, 2007.
- [10] A. Raftery. Hypothesis testing and model selection. In W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*, pages 163–188. Chapman and Hall, 1996.
- [11] Nestor V. Queipo, Raphael T. Haftka, Wei Shyy, Tushar Goel, Rajkumar Vaidyanathan, and P. Kevin Tucker. Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41:1–28, 2005.
- [12] A. O’Hagan. Bayesian analysis of computer code outputs: a tutorial. *Reliability Engineering and System Safety*, 91:1290–1300, 2006.
- [13] David Higdon. A primer on space-time modeling from a Bayesian perspective. In Barbel Finkenstadt, Leonhard Held, and Valerie Isham, editors, *Statistical methods for spatio-temporal systems*. CRC Press, 2007.
- [14] M. Frangos, Y. M Marzouk, K. Willcox, and B. van Bloemen Waanders. Surrogate and reduced-order modeling: A comparison of approaches for large-scale statistical inverse problems. In L. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. Mallick, Y. Marzouk, L. Tenorio, B. van Bloemen Waanders, and K. Willcox, editors, *Large-Scale Inverse Problems and Quantification of Uncertainty*. John Wiley & Sons, Ltd, Chichester, UK, 2010.

- [15] N. Wiener. The homogeneous chaos. *American Journal of Mathematics*, 60:897–936, 1938.
- [16] R.G. Ghanem and P.D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer Verlag, New York, 1991.
- [17] R. Ghanem. Probabilistic characterization of transport in heterogeneous media. *Computational Methods in Applied Mechanics and Engineering*, 158:199–220, 1998.
- [18] R.G. Ghanem, J.R. Red-Horse, and A. Sarkar. Modal properties of a space-frame with localized system uncertainties. In A. Kareem, A. Haldar, B.F. Spencer Jr., and E.A. Johnson, editors, *8th ASCE Specialty Conference of Probabilistic Mechanics and Structural Reliability*, number PMC200-269. ASCE, 2000.
- [19] D. Xiu and G.E. Karniadakis. Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos. *Computer Methods in Applied Mechanics and Engineering*, 191:4927–4948, 2002.
- [20] B.J. Debuschere, H.N. Najm, A. Matta, O.M. Knio, R.G. Ghanem, and O.P. Le Maître. Protein labeling reactions in electrochemical microchannel flow: Numerical simulation and uncertainty propagation. *Physics of Fluids*, 15(8):2238–2250, 2003.
- [21] O.P. Le Maître, R.G. Ghanem, O.M. Knio, and H.N. Najm. Uncertainty propagation using Wiener-Haar expansions. *Journal of Computational Physics*, 197(1):28–57, 2004.
- [22] X. Wan and G. E. Karniadakis. An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *Journal of Computational Physics*, 209:617–642, 2005.
- [23] K. Sargsyan, B. Debuschere, H. Najm, and O. Le Maître. Spectral representation and reduced order modeling of the dynamics of stochastic reaction networks via adaptive data partitioning. *SIAM Journal on Scientific Computing*, 31(6):4395–4421, 2010.
- [24] Robert B. Gramacy and Herbert K. H. Lee. Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103(483):1119–1130, 2008.
- [25] Robert B. Gramacy and Herbert K. H. Lee. Adaptive design and analysis of super-computer experiments. *Technometrics*, 51(2):130–145, 2009.
- [26] Youssef M. Marzouk, Habib N. Najm, and Larry A. Rahn. Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224:560–586, 2007.
- [27] X. Ma and N. Zabaras. An efficient Bayesian inference approach to inverse problems based on an adaptive sparse grid collocation method. *Inverse Problems*, 25:035013, 2009.

- [28] C. Lieberman, K. Willcox, and O. Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542, 2010.
- [29] P. G. Constantine, A. Doostan, Q. Wang, and G. Iaccarino. A surrogate accelerated bayesian inverse analysis of the HyShot II supersonic combustion data. In *Proceedings of the Summer Program 2010, Center for Turbulent Research*, 2010.
- [30] L. Mathelin, C. Desceliers, and M. Hussaini. Stochastic data assimilation of the random shallow water model loads with uncertain experimental measurements. *Computational Mechanics*, 47:603–616, 2011.
- [31] K. Sargsyan, C. Safta, B. Debusschere, and H. Najm. Multiparameter spectral representation of noise-induced competence in *Bacillus Subtilis*. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, PP(99):1, 2012.
- [32] S. A. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. *Soviet Mathematics Dokl.*, 4:240–243, 1963.
- [33] Knut Petras. Smolyak cubature of given polynomial degree with few nodes for increasing dimension. *Numerische Mathematik*, 93:729–753, 2003.
- [34] Knut Petras. Fast calculation of coefficients in the Smolyak algorithm. *Numerical Algorithms*, 26:93–109, 2001.
- [35] Knut Petras. On the Smolyak cubature error for analytic functions. *Advances in Computational Mathematics*, 12:71–93, 2000.
- [36] Thomas Gerstner and Michael Griebel. Numerical integration using sparse grids. *Numerical Algorithms*, 18:209–232, 1998.
- [37] Thomas Gerstner and Michael Griebel. Dimension adaptive tensor product quadrature. *Computing*, 71:2003, 2003.
- [38] A. Arsenlis et al. Enabling strain hardening simulations with dislocation dynamics. *Modelling and simulation in material science and engineering*, 15:553–595, 2007.
- [39] J. Marian and A. Caro. Moving dislocations in disordered alloys: Connecting continuum and discrete models with atomistic simulations. *Physical Review B*, 74(024113), 2006.
- [40] M. R. Gilbert, S. Queyreau, and J. Marian. Stress and temperature dependence of screw dislocations mobility in  $\alpha$ -Fe by molecular dynamics. *Physical Review B*, 84(174103), 2011.

## DISTRIBUTION:

1	Jaideep Ray, 08954	MS 9159
1	H. N. Najm, 08351	MS 9051
2	A. Arsenlis & S. Aubry	MS L-367, LLNL, Livermore, CA
1	M. Rhee,	MS L-129, LLNL, Livermore, CA
1	Technical Library, 09536 (electronic)	MS 0899



